Keeping it Real: Using Real-World Problems to Teach AI to Diverse Audiences

Nicole Sintov, Debarun Kar, Thanh Nguyen, Fei Fang,
Kevin Hoffman, Arnaud Lyet, and Milind Tambe

Abstract
In recent years, AI-based applications have increasingly been used in real-world domains. For example, game theory-based decision aids have been successfully deployed in various security settings to protect ports, airports, and wildlife. This paper describes our unique problem-to-project educational approach that used games rooted in real-world issues to teach AI concepts to diverse audiences. Specifically, our educational program began by presenting real-world security issues, and progressively introduced complex AI concepts using lectures, interactive exercises, and ultimately hands-on games to promote learning. We describe our experience in applying this approach to several audiences, including students of an urban public high school, university undergraduates, and security domain experts who protect wildlife. We evaluated our approach based on results from the games and participant surveys.

Introduction
Security is a global concern. A fundamental challenge of protecting critical infrastructure (e.g., ports, airports) as well as critical resources (e.g., fisheries, wildlife) arises from limited availability of security resources. Protecting all targets at all times is typically not realistic, and as a result, 100% security is not possible. Instead, security resources must be deployed intelligently, an endeavor in which artificial intelligence (AI) can play a major role.

Security Games: Using AI to Address Real-World Problems
In recent years, the field of security games, a subfield of AI, has drawn increasing attention from outside the AI community (Tambe 2011). In particular, game theory-based decision aids have been successfully deployed to protect critical infrastructure such as airports and ports (e.g., Pita et al. 2008), making real-world impacts and resulting in fundamental changes to security operations for various organizations. Security problems continue to evolve worldwide, creating new research challenges and practical applications for security games. Focusing on wildlife protection specifically, poaching represents the second largest threat to biodiversity after habitat destruction. This led to the development of green security games, a subfield of security games focused on protecting forests (Johnson et al. 2012), fisheries (Haskell et al. 2014) and wildlife (Fang et al. 2015; Kar et al. 2015). Although park rangers conduct patrols to combat poaching, security resources are often limited in vast conservation areas. Manually generating patrol schedules can require considerable effort from wildlife security staff, and such manual plans can be predictable, allowing poachers to exploit patrol schedules. Our security game-based solutions combine different AI subfields—including game theory, optimization, and machine learning—to help rangers automatically generate randomized patrol strategies that account for models of poachers' behaviors.

As a subfield of computational game theory, security games (Tambe 2011) model the strategic interaction between two players: a defender and an adversary. Security games take into account: (1) differences in the importance of targets; (2) the responses of attacker (e.g., poacher) behavior to the security posture; and (3) potential uncertainty over the types, capabilities, knowledge, and priorities of attackers. This problem can be cast as a game. As a brief example, a security game in the wildlife domain involves the following: the ranger allocates security resources (i.e., ranger patrol teams) to protect a set of critical targets of varying importance (Figure 1). Higher value "targets" may be portions of a protected area with higher biodiversity, larger numbers of animals, and/or protected species. The ranger deploys a mixed strategy, which optimizes over all possible configurations of allocating patrols across these targets, and is represented as a vector of probabilities of covering any given target. The poacher conducts surveillance on the ranger's strategy before selecting a target to attack, with the goal of maximizing payoff for any given defender strategy. The players' actions lead to different payoff values, and the defender's performance is evaluated by her/his *expected utility*. The defender's goal is to find the optimal strategy so as to maximize expected utility, knowing s/he faces an adaptive adversary who will respond to any deployed strategy.

Prior Work: Teaching with Projects, Problems, and Scaffolds
The potential applications of this work to various contexts have created the need to introduce the AI concepts underlying security games to individuals with limited AI backgrounds. These include not only students, but also audiences outside of traditional classroom settings. Given recent advances in green security games in particular, helping decision-makers and those who may consider using AI-based decision aids in the field to understand the underlying theoretical framework can aid in fostering adoption of these emerging technologies.

Teaching security games and related concepts such as probability, optimization, and agent-based modeling to those with limited AI backgrounds can be challenging. In traditional classroom settings, AI concepts have been made accessible to undergraduate students who enter with limited AI backgrounds (Stern and Sterling 1996; Parsons and Sklar 2004; Wollowski 2014). One method that has been effective in teaching AI in classrooms is the use of games. For instance, games have been used to teach robotics, (Wong et al. 2010), Pac-man has been used as a tool to teach various AI concepts (DeNero and Klein 2010), and in a game called CyberCIEGE, players build a virtual world while learning about AI issues involved in cyber security (Cone et al. 2007). However, no prior work describes effective methods for teaching AI to audiences beyond the classroom. Similarly, little evidence speaks to approaches for framing such games to teach and foster interest in AI.

We explored the possibility of using real-world problems to frame AI instruction. This approach is similar to project-based learning, an educational framework that aims to increase motivation for learning by engaging students in investigation (Blumenfield et al. 1991). Specifically, project-based learning involves presenting a problem that guides activities, and such activities culminate in a final product to answer the initial question. A meta-analysis of project-based learning studies conducted in real-world classrooms found that such an approach results in positive effect on application of general science knowledge, and although no immediate main effect on declarative knowledge (of underlying concepts, facts) was found, this increased over time (Dochy et al. 2003). Similar approaches have been applied to engineering curricula at the college level at several higher education institutions, and although no systematic evaluation results could be identified, qualitative feedback from students indicated that they evaluated the approach positively (Mills and Treagust 2003). Particularly relevant to AI, Gini et al. (1997) used a variety of robotics projects to teach robotics and other AI concepts at the college level; however, no evaluation results were provided. In sum, project-based learning has shown promise in piquing student interest and improving application of knowledge in general, but its effectiveness as a technique for teaching AI specifically, particularly outside of the classroom, has not yet been investigated.

Project-based learning also lends itself well to other instructional strategies. For instance, *scaffolding* is an instructional framework that can help learners achieve learning goals in an assisted, often step-wise manner (Wood et al. 1976). Broadly, in instructional scaffolding, teachers provide adjustable support for learners to enhance learning and promote progressive mastery of material by introducing new concepts and skills in a systematic manner (Pea 2004). Within a scaffolding framework, imaginary problems have been used to highlight the larger importance of learning goals, followed by project-based learning which can serve to structure specific learning activities and tasks (Barron et al. 1998). Such a "problem-to-project" scaffolding approach has previously been found to improve a variety of classroom learning outcomes (Barron et al. 1998). However, we could find no evidence speaking to its efficacy for teaching AI specifically, for teaching outside of traditional classroom settings.

Real-World Problem-to-Project Teaching Approach

Building on our prior paper (Sintov et al. 2016), this article describes our teaching approach, which used a problem-to-project scaffolding framework. The program began by presenting real-world wildlife security problems that painted a broad picture for why learning security game and other AI concepts is important. As detailed in the following sections, we then used project-based learning techniques along with instructional support to introduce progressively complex AI concepts underlying security games, including probabilistic reasoning, optimization, and agent-based modeling. We taught these concepts with a combination of lectures, interactive exercises, and hands-on games to help learners tie learning activities to the larger real-world security goals. We describe our experiences delivering this approach to several audiences, including: (1) students of an urban public high school; (2) undergraduate students at a large private university; and (3) law enforcement officers and rangers who protect wildlife in Indonesia. It is important to note that although the wildlife security problems used in our program are based on real-world data and input from security experts working in the field on such problems, they represent an abstracted version of the problems they aim to address.

We evaluate our teaching approach in several ways: (1) we show how high school and university participants played the games and the effectiveness of their strategies playing as defender, represented as *defender utility*, i.e., the defender's overall payoff obtained when peers play as poachers; (2) we show how security experts who typically generate defender strategies in the real world played as poachers against an AI algorithm-generated defender strategy; and (3) we assess participants' perceptions of our approach using surveys. We make recommendations for teaching similar topics to audiences with limited AI backgrounds.

Our Ranger vs. Poacher Games

We used interactive games as a key teaching tool in our program. Prior to playing these games, learners participated in lectures, discussions, and other learning activities. The games therefore provided learners an opportunity to apply their culmination of knowledge, while also allowing them to build on this knowledge through hands-on exercises that allowed trial and error testing of their ideas.

Adapting an online computer game used in another study of green security games (Kar et al. 2015), we created a board game. As background information, we describe the computer game here.

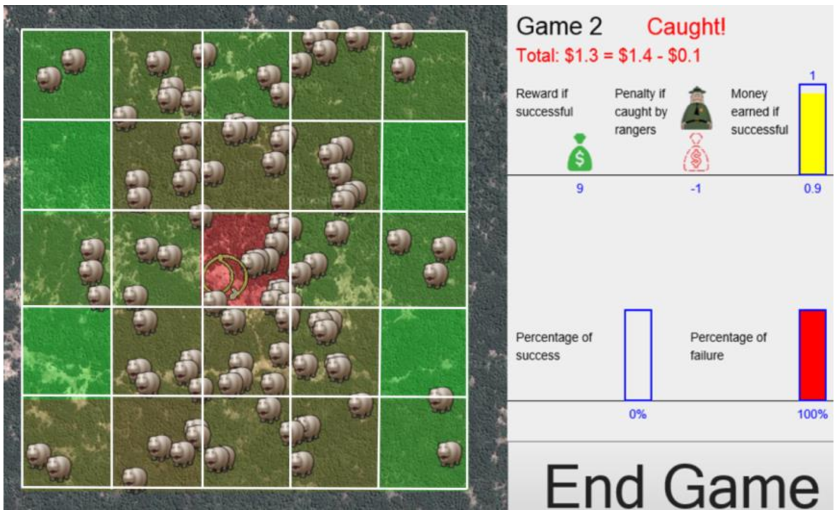

Figure 1. Online computer game interface (Kar et al. 2015).

Figure 2. Board game (Sintov et al., 2016).

The computer game was set in a protected wildlife park that was divided into a 5-by-5 grid, yielding 25 distinct cells, with hippopotami and rangers scattered throughout (Figure 1). In the computer game, participants played the role of a poacher whose objective was to hunt as many hippopotami as possible while avoiding detection by rangers. They were primed by reading a background story detailing the hardships of a poacher's life as well as the rewards of successful poaching. They could navigate throughout the park and select any cell to "attack". They were asked to consider three main criteria in deciding where to attack: *distance*, *animal density*, and *coverage probability*. Distance was the distance from the starting position to the attack location. This was incorporated to better approximate the real-world scenario whereby time and distance traveled to location a snare represents a cost to poachers. Animal density was represented by hippopotamus density, which varied across the park, but density in a particular region (cell) did not change within a given round. Coverage probability was represented by a heat-map overlaid on the park, indicating the likelihood of ranger presence at a given cell. Cells with higher coverage probabilities were shaded more red, whereas those with lower coverage probabilities were shaded more green. This represents the real-world surveillance situation in which poachers have knowledge of general patterns of ranger locations, but cannot always predict exact ranger presence. In the computer game, a computer algorithm determined ranger locations. A total of nine rangers were protecting the park, with each ranger protecting one cell. Therefore, only 9 of 25 distinct cells in the park were actually protected, and the color shaded heat-map showed players only the *probability* of ranger presence at each cell. Additionally, the following detailed information was available by clicking on any cell: monetary reward for poaching successfully, monetary penalty for getting caught, and chance of success/failure (figure 1). Players were successful in the game if they attacked a cell without a ranger and failed if a ranger was protecting their chosen cell.

For our teaching program, we adapted this computer game into a board game (figure 2) to provide learners the opportunity to play as both poachers and rangers. Board games used the same background story, 5-by-5 grid, defender resource allocation, and reward distribution as the computer game. However, instead of using the Queen Elizabeth National Park map and images of hippopotami to show the relative "value" of different areas, the board games used movable figures to represent animal distributions. This did not require any equipment such as computers, making it easily scalable to other settings. Also, in board games, learners could take turns playing the roles of ranger and poacher, whereas the computer game permitted play only as poachers. Additionally, the board games facilitated peer interaction. One group, acting as defender, was given a limited number of defense resources. They could place ranger figures anywhere on the board to generate the defender strategy. The other group, acting as adversary, placed the poacher figure anywhere on the board in deciding where to "poach" against their peers' defender strategy, representing poaching decision making. Finally, the board game did not provide immediate feedback on probabilities and rewards associated with success or failure (right-hand panel of Figure 1), better approximating the real-world situation in which this level of detail is typically unavailable. Therefore, the board games were leveraged as a more scalable and flexible learning tool. Although we aimed to abstract the real-world problem to the extent possible in this game, not all details of the real-world scenario could be included, and hence findings should be viewed in light of this limitation.

High School Students

An AI unit focused on computational game theory was delivered to a group of high school students as the last unit of study in a year-long engineering elective course. An underlying tenet of the unit was that AI concepts can be made accessible to anyone. We reasoned that AI concepts could also provide a great launching point into a discussion on computers in general; if students could understand quantitative decision making in this context, they could use it as a base of knowledge from which other computer science concepts could be more readily understood. Unit objectives included students gaining probabilistic reasoning skills, enhancing student interest in AI, and high levels of student satisfaction with the learning experience. The unit was developed during summer 2014, funded by the National Science Foundation's Research Experience for Teachers program called ACCESS 4Teachers, and was based on an undergraduate-level course at the University of Southern California entitled CS499: Artificial Intelligence and Science Fiction (Tambe et al. 2008).

Participants. The 30 students who participated in the elective course were juniors and seniors at an urban public charter high school located in Los Angeles. All students were of Hispanic or African-American origin, resided in South Los Angeles, and the majority qualified for free or reduced price lunch. Based on a subset of students who responded to our feedback survey, the mean age of the group was 18.1 years and roughly 36% were female.

Unit Structure. The unit began with a basic exploration of the nature of human intelligence and how machines (both fictional and real) have been made to mimic human nature. Consistent with a scaffolding framework, we next placed an emphasis on thinking about problems from a quantitative perspective, and considering how human-like qualities (emotions, risk-aversion, etc.) could be quantified to enable computers to act in an intelligent fashion. To illustrate these ideas, students read stories from Asimov's "I, Robot" and "Robot Visions", watched clips of the character Data from *Star Trek*, and debated the meaning of intelligence based on these stories. Students tended to be very engaged in these real-world examples, and were often perplexed when they saw how quantitatively-based decisions differed from their emotional ones. Next, to review the concept of probability, more pop-culture examples were given, ranging from an exploration of the California Lottery to the popular game show *Deal or No Deal?*, and the unit progressed to introducing the concept of expected value. The teacher delivered lectures, and skills were reinforced by students completing worksheets and applied problems. For many students, a full grasp of the material required a review of fraction operations, which were reviewed on an ad-hoc basis in small groups. The unit culminated with students integrating and applying their knowledge to play the games.

Final Project. The final project for the unit used the board game described above (figure 2). To reinforce the idea of quantitative decision-making, students were tasked with designing their own defender strategies using short formulas to allocate a limited number of ranger-hours to the 25 grid cells in the game. Students worked independently or in groups to complete the strategies (in this case, using Google spreadsheets). These spreadsheets were then used to generate twelve distinct games based on the student-designed defense strategies. Taking on the role of poacher, students played the games against defender strategies generated by other groups. Finally, students reviewed the results of their strategies, made adjustments, and presented their work to explain where initial strategies were particularly successful or unsuccessful.

Results of Games. Different groups employed different methods in designing defender strategies. While some chose to concentrate ranger patrols in areas dense with animals, often associated with a high probability of failure, others developed strategies in which the expected value for poachers was nearly zero. Figure 3 shows the expected defender utilities obtained by each of the 12 student groups based on the attacks conducted on the other teams' defender strategies. The team with the highest expected utility generated a strategy that not only considered the animal densities, but also the distance from the poacher's starting location, placing lower coverage in cells farther away from the poacher's starting location. The team with the lowest expected utility placed maximum coverage (100%) on the highest animal density cell and divided remaining resources (ranger-hours) uniformly across all remaining targets, ignoring important factors like animal distribution and distance.

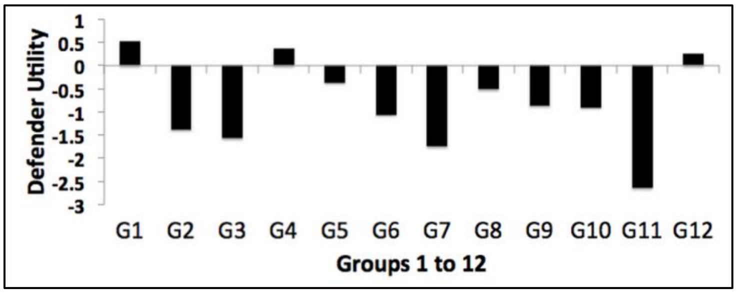

Figure 3. Defender utilities for high school student groups.

Feedback. A total of 14 (8 males, 5 females, 1 declining to state) out of 30 students responded to a survey that assessed their experiences in the unit. To inform unit objectives, questions assessed the unit's impact on learners' overall interest in AI, perceived educational value of the unit, and likelihood of recommending the unit to others. Responses were provided on Likert scales (e.g., ranging from 1=strongly disagree to 7=strongly agree). Open-ended questions assessed general likes and dislikes. More than 70% of respondents agreed (somewhat or more) that the activity increased their interest in AI, and 93% agreed (somewhat or more) that the activity was a valuable learning experience. Additionally, more than 90% responded that they would recommend the unit to other high school students. Open-ended responses indicated that respondents particularly enjoyed the interactive game aspect of the course. The least enjoyable aspects of the course were those that students perceived as tedious or repetitive, such as calculating probabilities.

University Students

AI and security games were introduced to a class of University of Southern California (USC) freshmen as a 2-week unit in their Freshman Academy course in fall 2015, which is an introductory engineering course aimed at introducing students to ongoing research at USC across various engineering disciplines. Similar to the high school course described in the previous section, the 2-week AI unit portion of the course was also designed based on a seminar at USC entitled CS499: Artificial Intelligence and Science Fiction. Aligning with unit objectives for the high schools students, objectives for the university students included honing probabilistic reasoning skills, enhancing student interest in AI, and high levels of student satisfaction with the learning experience.

Participants. The 30 students who took part in the AI unit were all USC freshmen majoring in engineering. Based on a subset of students who responded to our feedback survey, the mean age of the group was 18.2 years, and roughly 77% were female.

Unit Structure. Instructors opened the unit by introducing security as a global concern, and highlighted problems specific to in wildlife security. Step by step, the unit introduced more complex concepts, starting with basic concepts in AI and game theory. As part of our scaffolding framework, to teach the notion of payoffs in a game context, the classic Prisoner's dilemma problem in game theory was introduced. Discussion was facilitated around this topic to provide foundational understanding regarding payoffs in the games (i.e., animal densities and penalties). Faculty and CEOs of technical startups then facilitated discussion around the use of AI applications to solve real-world security problems, painting a picture of the various ways in which AI can influence day-to-day life. Similar to the high school students, the unit culminated with students integrating and applying their knowledge to play the games. Basic game-theoretic concepts such as *Maximin* were explained to help students to focus on subsets of information in decision-making: for instance, in the case of Maximin, when only information about the payoffs in the game is available, instructors aimed to help students design the most conservative strategy.

Final Project. Similar to the final project for high school students described above, following the lecture and discussion-based elements of the unit, students played the board game. Here, students first played as rangers. The class was divided into seven groups, each of which designed its own defender strategy on a game board. Some groups chose to allocate ranger coverage in proportion to the number of animals, whereas others placed highest coverage at the highest animal density region and uniformly everywhere else; some others developed strategies in which the expected value for poachers was nearly zero across all of the regions of the park. Each group's strategy was then shown to the other groups, who played the game in the poacher role against their peers' defender strategies.

*Results of Games***.** The resulting defender utilities for each university student group playing the board games is shown in figure 4. The team with the lowest defender utility (G3) placed very low coverage (< 0.40) in the highest animal density cell and as expected, all the other teams attacked that cell. Similar results were obtained for group G7 which placed a coverage of 0.50 on the highest animal density target. Other teams performed reasonably well but none performed better than Maximin, which would have resulted in a defender expected utility of -2.17.

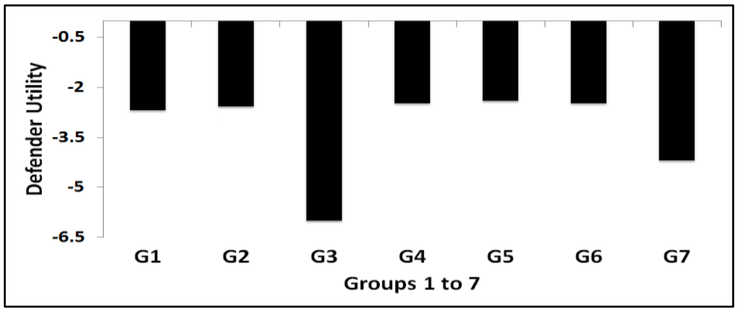

Figure 4. Defender utilities for university student groups.

Comparing the university students' results with those of the high school students (figure 3), visual inspection suggests that overall, the high school students outperformed the university students. This may be due to the limited time allotted to the university students. On the other hand, the high school students had more time, which allowed more discussion and scrutinizing over decisions. However, the small sample sizes prohibited statistical comparison, so this observation should be interpreted with caution.

Feedback. A total of 24 (7 males and 17 females) out of 30 university students responded to a survey that assessed their experiences in the unit. Questions mirrored those administered to the high school students to address unit objectives. More than 69% of respondents indicated that the activity increased their interest in AI at least somewhat, and more than 80% agreed (somewhat or more) that the activity was a valuable learning experience. Additionally, more than 65% responded that they would recommend the activity to peers. Qualitative data suggested that respondents particularly enjoyed the interactive aspects of the unit. The least enjoyable aspects of the unit were cases in which students in a team couldn't agree on a particular strategy.

Security Experts

A 3-day workshop was developed in collaboration with the World Wildlife Fund (WWF) to demonstrate the value of AI-based solutions for security to security experts who protect wildlife. The workshop was held in Bandar Lampung, Sumatra, Indonesia in May 2015. A game-theory based decision-aid called PAWS (Yang et al. 2014; Protection Assistant for Wildlife Security) was developed, in part, based on a study of green security games for the purpose of protecting wildlife from poaching. We sought to teach how AI systems like PAWS fed with partial information can generate patrol strategies that can perform relative to strategies created by field experts with extensive knowledge of the system. Hence, diverging a bit from the classroom-based units described above, our objectives for this third audience included promoting participant adoption of AI-based software (i.e., PAWS), sharpening participants' probabilistic reasoning skills especially in the poacher role, and participant satisfaction with the learning experience.

Participants. A total of 28 participants (26 males and 2 females) attended the workshop. They represented the five primary groups (either government or NGO) involved in protecting wildlife in Bukit Barisan Selatan and Tesso Nilo national parks on Sumatra: the Indonesian National Park Service, WWF, Wildlife Conservation Society, Indonesian Rhino Foundation, and Prosecution Officers from the Court. The majority of these individuals were rangers with a great deal of domain expertise in wildlife crime and protection who directly conduct field patrols over conservation areas; the prosecutors report cases to lawyers and judges who can open official investigations to prosecute wildlife crime. The mean age of the sample was 35.0 years (SD = 7.5), and mean years of formal schooling was 14.0 (SD = 3.1). Approximately 60% of respondents identified their job sector as wildlife/national park protection, 20% as nonprofit/NGO, and 20% as law enforcement, and overall they had an average of 9.6 years of experience working in wildlife protection (SD = 6.1).

Participants were native speakers of Bahasa Indonesia. The instructors delivered the workshop in English and interpreters translated all the material between instructors and participants throughout the 3-day course. All written materials were made available in both English and Bahasa Indonesia.

Unit Structure. We began by introducing basic examples and theoretical foundations relevant to agent-based modeling, game theory, and security games via lectures. Building on and integrating this knowledge as part of our scaffolding framework, we next presented applications that leverage multiple AI techniques. Learners also discussed in groups various challenges faced in wildlife protection and solutions for those challenges, including AI-based solutions. They played the computer-based game as poachers. On the last day, they had the opportunity to integrate and apply their knowledge in playing the board game as poachers and rangers. They also reflected on their results and shared ideas for improving patrolling effectiveness. These interactive exercises provided learners with a new lens for understanding poachers' behaviors and limits of manual patrolling strategies, as well as introducing the methodology and advantages of game-theoretic solutions.

Security Game Tutorials. On the first day of the workshop, we introduced security game examples from several domains, beginning with a basic security game. We explained how the defender could optimally conduct patrols over targets and how attackers may respond against that strategy. We next covered (1) real-world applications of security games for protecting critical infrastructure and (2) challenges in wildlife protection and the application of security games to this domain specifically, reviewing how adversaries' behaviors are modeled and how to optimize patrolling strategies through allocating limited security resources. Finally, we presented our PAWS software—which was built based on security game models for addressing wildlife protection problems—describing data inputs (e.g., animal density, poaching data) and what outputs are generated (e.g., models of poachers' behaviors and suggested patrolling routes for rangers). We described how similar approaches had been previously successfully used in the wildlife domain, and how they could be used on Sumatra with PAWS.

Discussion Sessions. Participants engaged in several discussion sessions on challenges in wildlife protection including resources (i.e., factors that motivate people to enter protected areas), illegal activities (i.e., types of illegal activities in conservation areas), and wildlife protection (i.e., improving security approaches). In small groups, they exchanged knowledge about these topics and generated potential solutions; each group then presented their conclusions to the rest of the groups. We encouraged groups to develop solutions and provide feedback that could be conceptualized in a game-theoretic manner and potentially incorporated into AI software.

Games. Participants played the board game as poachers and rangers. For this activity, participants were divided into two groups. Each group took turns playing as rangers (who created patrol strategies) and poachers (who decided where to poach in games generated by the other team), and each defender strategy was played only once.

Given the large amount of time for the workshop relative to the classroom-based units, in addition to board games, every participant played five rounds of the computer-based games as poachers. After each round, the poacher behavior models were updated based on participants' responses, and each subsequent game used a defender strategy created using these updated models. On the final day, we presented the game results, i.e., the defender utilities based on poachers' decisions in the online games.
By playing these games in a repeated fashion, the participants developed a better understanding of how poachers may react to rangers' strategies over time, and of the weaknesses of various defender strategies. They also learned how AI software such as PAWS can make optimal decisions based on models of players' behaviors, and how such decisions can adapt and improve over time as more data are collected.

Figure 5. Participants played computer games (left) and board games (right).

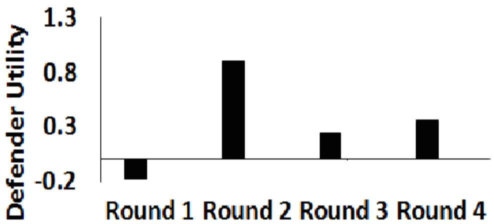

Figure 6. PAWS algorithm defender utilities against security experts playing as poachers in computer game.

Results of Games. Each defender strategy in the board games was played only once, so the results based on single data points may not be reliable. In light of this and the fact that only the security experts played the computer game (whereas high school and university students did not), we highlight the results of the computer games here. Figure 6 shows the defender utilities obtained by deploying AI-based defender strategies (i.e., PAWS) over several rounds against security experts playing as poachers. In the figure, lower values on the y-axis indicate better participant performance, and worse performance by PAWS. We observe that PAWS's performance begins low, initially increases, then declines. This suggests some improvement and learning over time among the security experts, providing modest support of our unit objective of improving probabilistic reasoning skills.

Feedback. A survey administered at the end of the unit addressed unit objectives. Approximately 79% of learners were at least somewhat willing to adopt PAWS (mean = 5.9 on a scale of 1 [strongly disagree] – 7 [strongly agree], SD = 1.1). Additionally, open-ended responses also largely supported the purpose of the PAWS software. Roughly half of respondents (n=15) commented that PAWS could optimize patrols and would make the job of patrolling easier. When asked about software limitations, respondents recommended increasing complexity of models, including approximately one-third (n=10) of respondents suggesting that dynamic animal distribution models be added. This latter point highlights the challenges faced by AI researchers in accurately representing all details of a given real-world scenario.

Regarding satisfaction with the unit, 86% rating the learning experience as at least somewhat useful (mean = 5.7 on a scale of 1 [completely useless] – 7 [extremely useful], SD = 0.8), and more than 96% of respondents rated it as at least somewhat important (mean = 6.0 on a scale of 1 [extremely unimportant] – 7 [extremely important], SD = 1.0). Additionally, more than 86% of respondents reported that they were at least somewhat likely to recommend it to peers (mean=6.04 on a scale of 1 [strongly disagree] – 7 [strongly agree], SD = 1.04).

Conclusion

This paper describes our unique approach that used a real-world problem-to-project scaffolding framework to teach game-theoretic concepts to several audiences. Learners included students at an urban public high school, university undergraduate students, and law enforcement officers and park rangers who protect wildlife in Indonesia. Our instructional units began by presenting real-world problems in wildlife and other security domains that painted the broad picture for why learning security game concepts is important. Throughout the learning units, techniques from project-based learning along with instructional support were used to progressively introduce complex AI concepts and help learners tie learning activities to the larger real-world security goals. Games were a key learning tool in our approach. Members of all three audiences played the role of two different actors as part of our games: (1) playing as rangers, they generated defender strategies to protect against poachers' attacks; (2) as poachers, they attempted to outsmart the defender strategies generated by their peers to earn the highest possible rewards. This approach not only gave learners valuable hands-on experience with complex AI concepts, but also in the development of real-world applications for security.

Participant feedback was consistently positive, with the majority of participants from all three audiences rating the learning experiences as useful, and indicated they would recommend the unit to others. In addition, the majority of the high school and university students reported that it increased their interest in AI. These results are particularly notable given the wide age range and cultural backgrounds of participants, and suggests that our approach was broadly accessible and engaging.

The benefits of problem-to-project scaffolding in our work included having a unifying theme, enabling both learners and instructors to tie concepts and tasks back to a central idea. A related benefit was that the theme afforded the ability to enhance relevance, especially for audiences outside of the classroom, highlighting the value of our approach for teaching concepts in AI to learners representing a broad swath of age, gender, and cultural backgrounds. Additionally, learners showed relatively high levels of engagement, as evidenced by survey findings and qualitative feedback. Being able to “meet students where they were at”, and serve different learning styles using a variety of activities, was another benefit of the broader scaffolding approach. Drawbacks of our scaffolding framework included somewhat of a narrow focus: by beginning with a particular security problem, a circumscribed range of AI topics “fit” with the unit. This also led to development of instructional materials specific to the problem.

These findings suggest that our approach may potentially be applied successfully to additional audiences. For instance, our approach can help address educators’ needs, as a recent study found that many are seeking more cutting-edge classroom materials (Wollowski et al. 2015). A possible additional target group is security organizations that could benefit from security applications based on AI and game theory. Enhancing decision-makers’ and field officers’ understanding of the theory on which these applications are based could foster the adoption of emerging AI-based decision aids.

Limitations and Future Directions

Our findings should be viewed in light of several limitations. First, it is important to note that although the wildlife security problems used in our program were based on real-world data and input from security experts working in the field on such problems, they represent an abstracted version of the problems they aim to address. In addition, it is unclear how our approach would generalize to teaching AI topics beyond game theory and security games. In the future, we plan to adapt our activities to focus on other AI topics. We also plan to incorporate new activities; for instance, an activity for learners to analyze defender strategies in depth, bringing to light subtle human biases that may impact initial strategies, thereby highlighting benefits of AI agents (compared to humans) in decision-making.

Acknowledgements

This work was supported by MURI Grant W911NF-11-1-03 and World Wildlife Fund PAWS Anti-Poaching Grant.

References

Barron, B. J.; Schwartz, D.L.; Vye, N. J.; Moore, A.; Petrosino, A.; Zech, L.; and Bransford, J.D. 1998. Doing with Understanding: Lessons from Research on Problem- and Project-Based Learning. *The Journal of the Learning Sciences* 7(3): 271-331.

Blumenfeld, P.C.; Soloway, E.; Marx, R.W.; Krajcik, J.S.; Guzdial, M.; and Palincsar, A., 1991. Motivating Project-based Learning: Sustaining the Doing, Supporting the Learning. *Educational Psychologist* 26(3): 369-398.

Cone, B.D.; Irvine, C.E.; Thompson, M.F. and Nguyen, T.D. 2007. A Video Game for Cyber Security Training and Awareness. *Computers and Security* 26(1): 63-72.

DeNero, J., and Klein, D. 2010. Teaching Introductory Artificial Intelligence with Pac-Man. In Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence. 1885-1889. Palo Alto: AAAI Press.

Dochy, F.; Segers, M.; Van den Bossche, P; and Gijbels, D. 2003. Effects of Problem-Based Learning: A Meta-Analysis. *Learning and Instruction* 13(5): 533-568.

Fang, F.; Stone, P.; and Tambe, M. 2015. When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing. In Proceedings of the 24th International Joint Conference on Artificial Intelligence, 2589-2595. Buenos Aires: International Joint Conferences on Artificial Intelligence.

Gini, M.; Hougen, D.; Ozerkovsky, A.; Rybski, P.; and Schmalz, B. 1997. Interacting with the Real World: A Way of Teaching Artificial Intelligence Concepts.

Haskell, W. B.; Kar, D.; Fang, F.; Tambe, M.; Cheung, S.; and Denicola, L. E. 2014. Robust Protection of Fisheries with COmPASS. In Proceedings of the Twenty-Sixth Annual Conference on Innovative Applications of Artificial Intelligence, 2978-2983. Palo Alto: AAAI Press.

Johnson, M. P.; Fang, F.; and Tambe, M. 2012. Patrol Strategies to Maximize Pristine Forest Area. In Proceedings of 26th Annual Conference on Innovative Applications of Artificial Intelligence, 295-301. Palo Alto: AAAI Press.

Kar, D.; Fang, F.; Fave, F. D.; Sintov, N.; and Tambe, M. 2015. A Game of Thrones: When Human Behavior Models Compete in Repeated Stackelberg Security Games. In Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, 1381-1390. Richland: International Foundation for Autonomous Agents and Multiagent Systems.

Mills, J.E. and Treagust, D.F. 2003. Engineering Education—Is Problem-Based or Project-Based Learning the Answer. *Australasian Journal of Engineering Education* 3(2): 2-16.

Parsons, S.; and Sklar, E. 2004. Teaching AI using LEGO Mindstorms. In Proceedings of American Association for Artificial Intelligence Spring Symposium, 43-39. Palo Alto: AAAI Press.

Pea, R.D. 2004. The Social and Technological Dimensions of Scaffolding and Related Theoretical Concepts for Learning, Education, and Human Activity. *The Journal of the Learning Sciences* 13(3): 423-451.

Pita, J.; Jain, M.; Western, C.; Portway, C.; Tambe, M.; Or-donez, F.; Kraus, S.; and Paruchuri, P. 2008. Deployed AR-MOR protection: The Application of a Game Theroetic Model for Security at the Los Angeles International Airport. In Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent systems: industrial track, 125-132. Richland: International Foundation for Autonomous Agents and Multiagent Systems.

Sintov, N., Kar, D., Nguyen, T., Fang, F., Hoffman, K., Lyet, A., & Tambe, M. (2016). From the Lab to the Classroom and Beyond: Extending a Game-Based Research Platform for Teaching AI to Diverse Audiences. Proceedings of the *AAAI 6th Symposium on Educational Advances in Artificial Intelligence.*

Stern, L. and Sterling, L. 1997. Teaching AI Algorithms using Animations Reinforced by Interactive Exercises. In Proceedings of the ACM SIGCSE 2nd Australasian Conference on Computer Science, 231:239. Melbourne: ACSE.

Tambe, M.; Balsamo, A.; and Bowring, E. 2008. Using Science Fiction in Teaching Artificial Intelligence. In Proceedings of American Association for Artificial Intelligence Spring Symposium, 86–91. Palo Alto: AAAI Press.

Tambe, M. 2011. Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned. In Communication, Control, and Computing 50th Annual Allerton Conference, 1822-1829. New York: Institute of Electrical and Electronics Engineers.

Wollowski, M. 2014. Teaching with Watson. In Proceedings of the Fifth Symposium on Educational Advances in Artificial Intelligence, 3044– 3049. Palo Alto: AAAI Press.

Wollowski, M.; Selkowitz, R; Brown, L.E.; Goel, A.; Luger, G.; Marshall, J.; Neel, A.; Neller, T.; and Norvig, P. 2015. A Survey of Current Practice and Teaching of AI. Proceedings of the *AAAI 6th Symposium on Educational Advances in Artificial Intelligence.*

Wong, D.; Zink, R.; and Koenig, S. 2010. Teaching Artificial Intelligence and Robotics via Games. In Proceedings of the Symposium on Educational Advances in Artificial Intelligence, 1917-1918. Palo Alto: AAAI Press.

Wood, D.; Bruner, J.S.; and Ross, G. 1976. The Role of Tutoring in Problem Solving. *Journal of Child Psychology and Psychiatry* 17(2): 89-100.

Yang, R.; Ford, B.; Tambe, M.; and Lemieux, A. 2014. Adaptive Resource Allocation for Wildlife Protection Against Illegal Poachers. In Proceedings of the 2014 International Conference on Autonomous Agents and Multiagent Systems, 453-460. Richland: International Foundation for Autonomous Agents and Multiagent Systems.